

Estimating Cross-Platform Chess Rating Mappings with Modal Regression

Elo+Chess Research Notes

May 31, 2026

Abstract

Players often ask how ratings on Lichess.org compare with ratings on Chess.com. The question is practically important for Elo+Chess because its benchmark curves are built from a very large stratified sample taken from the complete Lichess.org monthly game histories, while users may bring histories from either site. This note describes the current method used to estimate game-type-specific mappings from Lichess ratings to Chess.com ratings. The method uses same-username candidate matches, exact or near-exact time-control filters, minimum game-history requirements, and a modal regression procedure designed to reduce bias from false same-username matches. The current current mapping estimate is:

$$\widehat{R}_{\text{Chess.com}} = \alpha + \beta R_{\text{Lichess}}.$$

The estimates are preliminary and will be updated as the matched sample grows, The May 31 update materially increases the longer-than-10-minute rapid matched sample and updates that curve accordingly.

Current Mapping Estimate

For readers who want the result first, the current fitted equations are:

$$\begin{aligned}\widehat{R}_{\text{Chess.com,bullet}} &= -531.71 + 0.9871R_{\text{Lichess}}, \\ \widehat{R}_{\text{Chess.com,blitz}} &= -551.54 + 1.0853R_{\text{Lichess}}, \\ \widehat{R}_{\text{Chess.com,10minrapid}} &= -495.04 + 1.0741R_{\text{Lichess}}, \\ \widehat{R}_{\text{Chess.com,>10rapid}} &= -369.02 + 0.9374R_{\text{Lichess}}.\end{aligned}$$

These lines are fitted over Lichess ratings 650–1600 and should not be extrapolated as universal conversions across all skill levels.

1 Objective

The goal is to map a rating on Lichess.org to the comparable rating scale on Chess.com for the same broad time-control category. We estimate a separate mapping for each category used by Elo+Chess:

- 1+0 bullet,
- 3+0 blitz,

- 10+0 rapid,
- rapid games longer than 10 minutes.

The fitted line is used only as a scale conversion layer. It does not imply that the two rating pools have identical rating mechanics, player pools, or game-selection behavior.

2 Data Provenance

The Lichess side of the analysis is built from the complete public Lichess standard rated game archives from January 2025 through March 2026. This matters because the Lichess sample is not a convenience scrape or a narrow sample of visible profiles: it is constructed from the full monthly game-history files released by Lichess for research and public use.

Lichess publishes these database exports on <https://database.lichess.org/>. The database page states that the exports are released under a “CC0 license” and explicitly permits users to “Use them for research” and “publication”.¹ This open-data posture is central to the design of the benchmark side of Elo+Chess.

On the Chess.com side, candidate same-username accounts are checked only for the monthly archives and game categories needed to estimate the mapping. We place special emphasis on matching game histories within the same calendar month whenever possible. Calendar-month matching reduces the risk that a player’s rating changed substantially between the Lichess and Chess.com measurements.

3 Candidate Match Construction

The main challenge is that neither platform exposes a universal cross-platform player identifier. We therefore begin with same-username candidates: if a username appears on both Lichess.org and Chess.com, that pair is treated as a candidate match. This is intentionally noisy. Some same-name accounts are the same person, while others are unrelated players who happen to use the same handle.

To reduce avoidable noise, the current pipeline applies game-type-specific filters:

1. Lichess users are sampled from stratified game-history datasets by rating bucket and time-control type.
2. Chess.com candidate accounts are checked for the corresponding time-control category.
3. Candidate pairs are retained only when both sides have enough relevant game history.
4. For bullet, blitz, and 10-minute rapid, the minimum retained history is 5 games per side.
5. For longer-than-10-minute rapid, the minimum retained history is 10 games per side.
6. The regression fit reported here is restricted to Lichess ratings from 650 to 1600, matching the rating region most relevant to the main Elo+Chess benchmark reports.

¹<https://database.lichess.org/>, accessed May 29, 2026. The quoted phrases are from the Lichess database page.

The current sample emphasizes same-calendar-month same-username matches. When a candidate account has sufficient games in the target time-control category in the same month as the Lichess source data, that matched monthly rating is preferred for fitting. Older cross-platform match data are retained where they increase sample support, but the methodological target is a game-type-specific same-month comparison.

4 Why Mean or Median Can Be Biased

For a fixed Lichess rating bucket, the same-username candidate set is a mixture:

$$f(y | x) = \pi f_{\text{true}}(y | x) + (1 - \pi) f_{\text{false}}(y),$$

where x is the Lichess rating, y is the Chess.com rating, f_{true} is the distribution of Chess.com ratings for true cross-platform matches, and f_{false} is the distribution generated by unrelated Chess.com accounts with the same username.

The false-match component is not centered on the correct Chess.com equivalent for the Lichess bucket. It tends to resemble the broad Chess.com population sampled by usernames and activity filters. As a result, the conditional mean is pulled toward the global active-player distribution:

$$E[y | x] = \pi E_{\text{true}}[y | x] + (1 - \pi) E_{\text{false}}[y].$$

If $\pi < 1$, this is not generally equal to the desired matched-player expectation. The same issue affects the median whenever the false-match component contributes enough mass to move the 50th percentile.

The direction of bias depends on the bucket. In low Lichess buckets, false matches can pull the mean upward if random active Chess.com accounts are stronger than the true matched users. In high buckets, the same contamination can pull the mean downward. This compression effect is exactly what we want to avoid when building a scale conversion.

5 Modal Motivation

The modal approach is based on a weaker assumption than the mean-based approach. We do not require every same-username pair to be correct. Instead, we assume that within a sufficiently narrow Lichess rating bucket, true matches form the most concentrated local cluster in Chess.com rating space. False matches may be numerous, but they should be more diffusely distributed because they are not tied to the Lichess rating bucket.

For each Lichess bucket, we estimate the mode of the Chess.com rating distribution using a kernel density estimate. The modal point is then treated as the most plausible Chess.com equivalent for that Lichess bucket. This makes the procedure less sensitive to long tails and less sensitive to mismatched same-name accounts sampled from the broader Chess.com population.

This is not a claim that the mode is perfect. It is a pragmatic robustness choice for a noisy entity-resolution problem. The method works best when true matches remain a plurality near the correct cross-platform rating and false matches are not themselves concentrated at a misleading value.

6 Fitting Procedure

The current production fit uses the following procedure separately for each game type:

1. Retain candidate same-username pairs satisfying the game-count thresholds.
2. Restrict the fit to $650 \leq R_{\text{Lichess}} \leq 1600$.
3. Partition Lichess ratings into 100-point buckets.
4. For each bucket with at least 30 users, estimate the Chess.com modal rating with a Gaussian-kernel density estimate using bandwidth 75 rating points.
5. Represent each bucket by its midpoint and estimated Chess.com mode.
6. Fit an ordinary least-squares line through the bucket-level modal points:

$$\widehat{R}_{\text{Chess.com}} = \alpha + \beta R_{\text{Lichess}}.$$

Only bucket-level modal points enter the final linear regression. Individual same-username pairs influence the curve through their contribution to the within-bucket density estimate.

7 Current Estimates

Table 1 reports the current fitted constants and slopes. Relative to the May 29, 2026 fit, the bullet, blitz, and 10+0 rapid equations changed only negligibly. The substantive update is the longer-than-10-minute rapid curve, where the same-username support increased from 750 users to 2,097 users after the latest targeted monthly matching pass.

Game type	Total users	Baseline	Same-month	In range	α	β
1+0 bullet	2,999	1,512	1,525	1,845	-531.71	0.9871
3+0 blitz	3,101	1,331	1,795	2,076	-551.54	1.0853
10+0 rapid	3,084	1,547	1,575	2,520	-495.04	1.0741
> 10 rapid	2,097	463	1,677	1,649	-369.02	0.9374

Table 1: Current modal linear fits over Lichess ratings 650–1600. “In range” counts users retained inside the fitted Lichess interval. Baseline rows come from the existing exact-control cross-platform table; same-month rows come from the March 2026 monthly same-username sample.

The longer-than-10-minute rapid estimate now has substantially stronger support than in the May 29 fit, although it remains somewhat less dense than the other three categories in the fitted range.

8 Interpretation

The fitted equations are best interpreted as local conversions for beginner to early-advanced users, not as universal rating identities. For example, the 10-minute rapid estimate is

$$\widehat{R}_{\text{Chess.com}} = -495.04 + 1.0741 R_{\text{Lichess}}.$$

At a Lichess rating of 1200, this gives a Chess.com equivalent of about 794. At a Lichess rating of 1500, it gives about 1116.

The estimates should not be extrapolated far outside the fitted interval. At very high ratings, player-pool composition, account behavior, and platform rating mechanics can differ in ways that are not captured by a single linear mapping fitted to the lower and middle rating range.

Cross-platform same-username modal rating fits

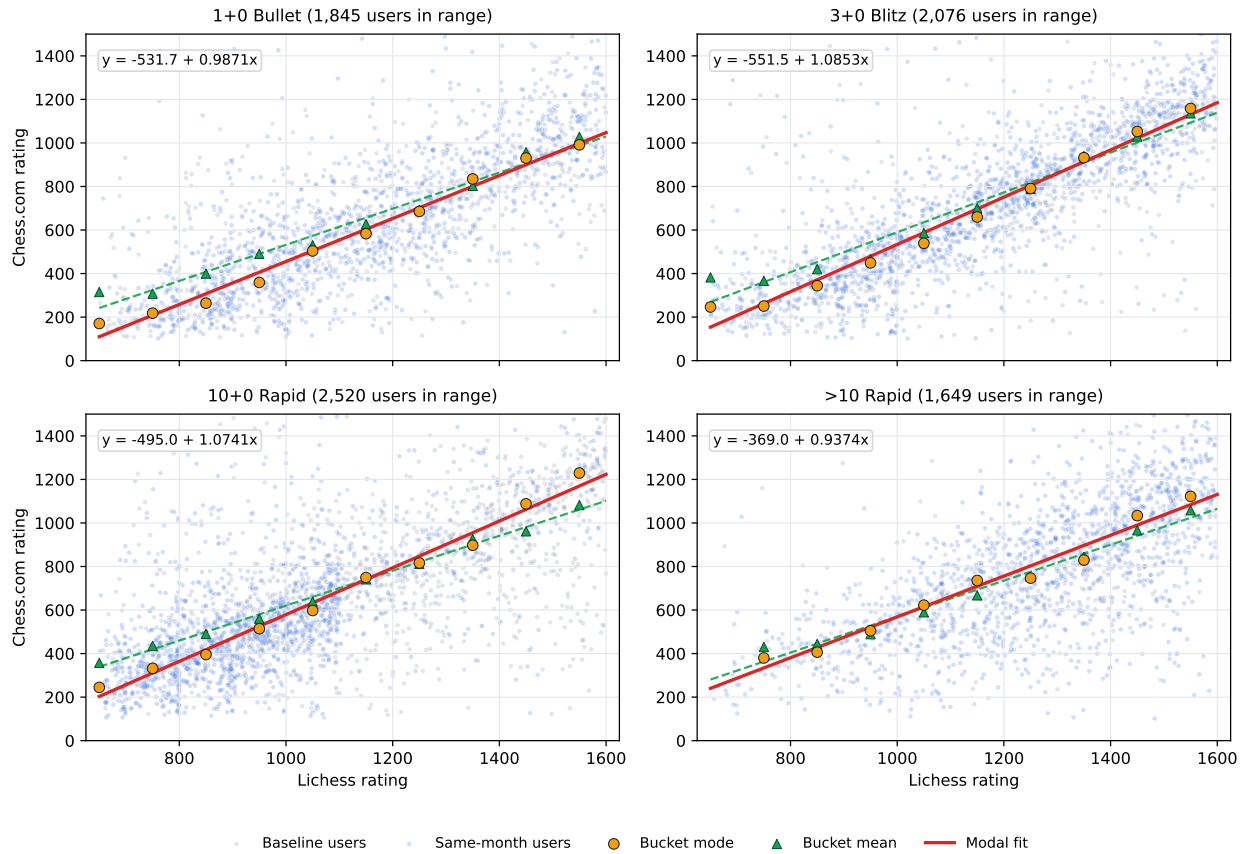


Figure 1: Same-username candidate matches and fitted mapping curves. Raw points are retained candidate users after game-history filters. Orange points are the within-bucket modal Chess.com ratings used for the production fit. The red line is the modal regression line. Green triangles and the dashed green line show the bucket-mean alternative for comparison; these are included to illustrate how mean-based summaries can be pulled by mismatched same-name accounts and tail observations.

9 Limitations and Update Plan

The current procedure is intentionally conservative, but several limitations remain:

- Same-username matching is noisy and does not prove identity.
- The true-match fraction may vary by rating bucket and game type.
- Chess.com and Lichess ratings are not always measured on the exact same day for all candidate pairs, although the same-month sample reduces this issue.
- The longer-than-10-minute rapid sample has improved materially, but remains less dense than the bullet, blitz, and 10-minute rapid samples.
- The linear model is a practical approximation over 650–1600, not a structural model of either rating system.

The next planned improvement is to continue sampling cautiously where it is useful, subject to conservative API usage and platform terms. Once the matched sample is larger and more balanced across rating buckets, the modal fits should be re-estimated and the production constants updated again.

10 Maintenance and Updates

This document reflects the most recent Elo+Chess cross-platform rating mapping update as of May 31, 2026. The estimates will be revised as additional same-calendar-month matched users are collected.

The maintained public version of this report is available at:

<https://www.elopluschess.com/static/docs/cross-platform-elo-mapping-modal-method.pdf>